



Microsoft®

**Windows.net™**

Server Family

# **File Server Part 1: Storage Management**

**Steve Olsson**

**Program Manager  
Core OS**

**Microsoft Corporation**

# **What's New**

- **Volume Shadow Copies**
- **Backup Infrastructure**
- **Automated System Recovery**
- **Command Line Scripting Tools**
- **Virtual Disk Service**

# **Volume Management:**

## **Basic vs Dynamic Volumes**

- **Which is right for your application**

- **Basic**

- **Supports growing Volumes on-line with Diskpart**
- **Best choice for using hardware Shadow Copies**

- **Dynamic**

- **Supports software fault tolerant configurations**
  - **Mirroring**
  - **Striping with parity**
- **Supports growing Volumes on-line with Diskpart or UI**



# Shadow Copies

- A Shadow Copy of a volume or group of volumes at a point in time
- An apparent or actual new volume will be created as hidden and read only
- A user or application may utilize this volume for any purpose
- Involves applications/services and file system to ensure data on disk is “consistent”

# Shadow Copy Scenarios

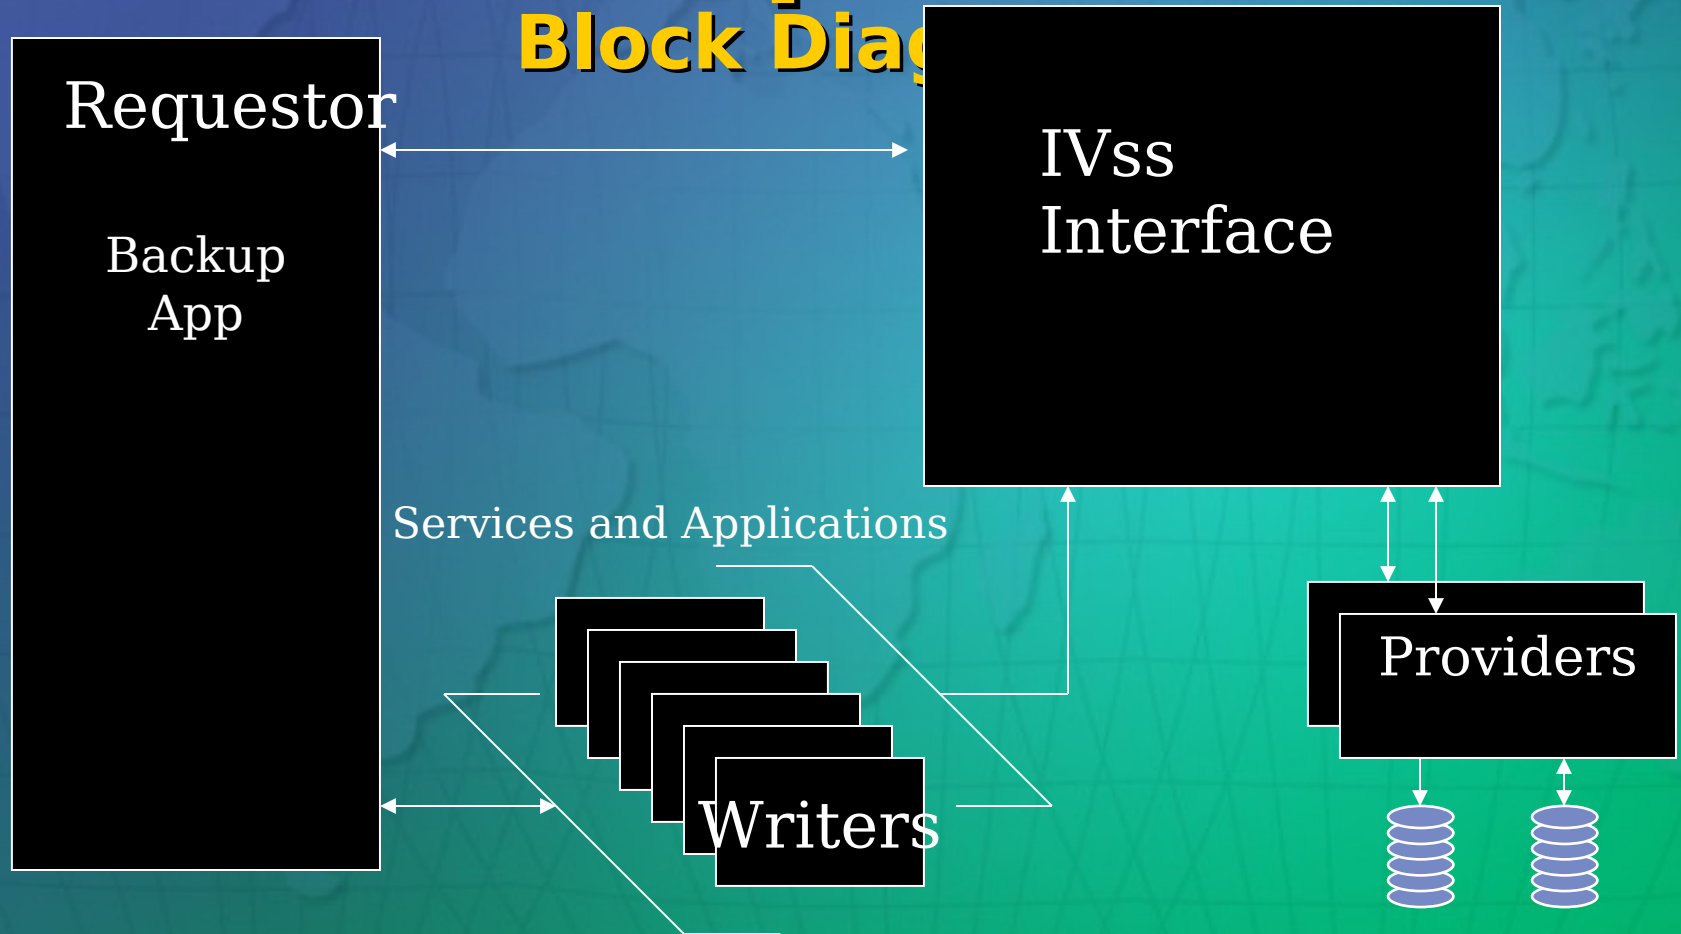
- **Open File Backup**
- **No more “Backup Window”**
- **Future Apps register “how to be backed up and restored”**
- **“Data Freighting”: Clone volumes and move to another host on a SAN**
- **Also used by “Shadow Copy Restore” covered in FP302 session**

# **Infrastructure Goals**

- **Application synchronization across volumes**
  - **Freeze/Thaw APIs for 'Trained Apps'**
  - **'Crash-consistency' for Legacy apps**
    - **All snapped at a point in time**
- **Discovery and enumeration of Shadow Copies**
  - **Even on a SAN where pieces have 'moved'**
- **Enable IHVs and ISVs to plug in their Shadow Copy engines**

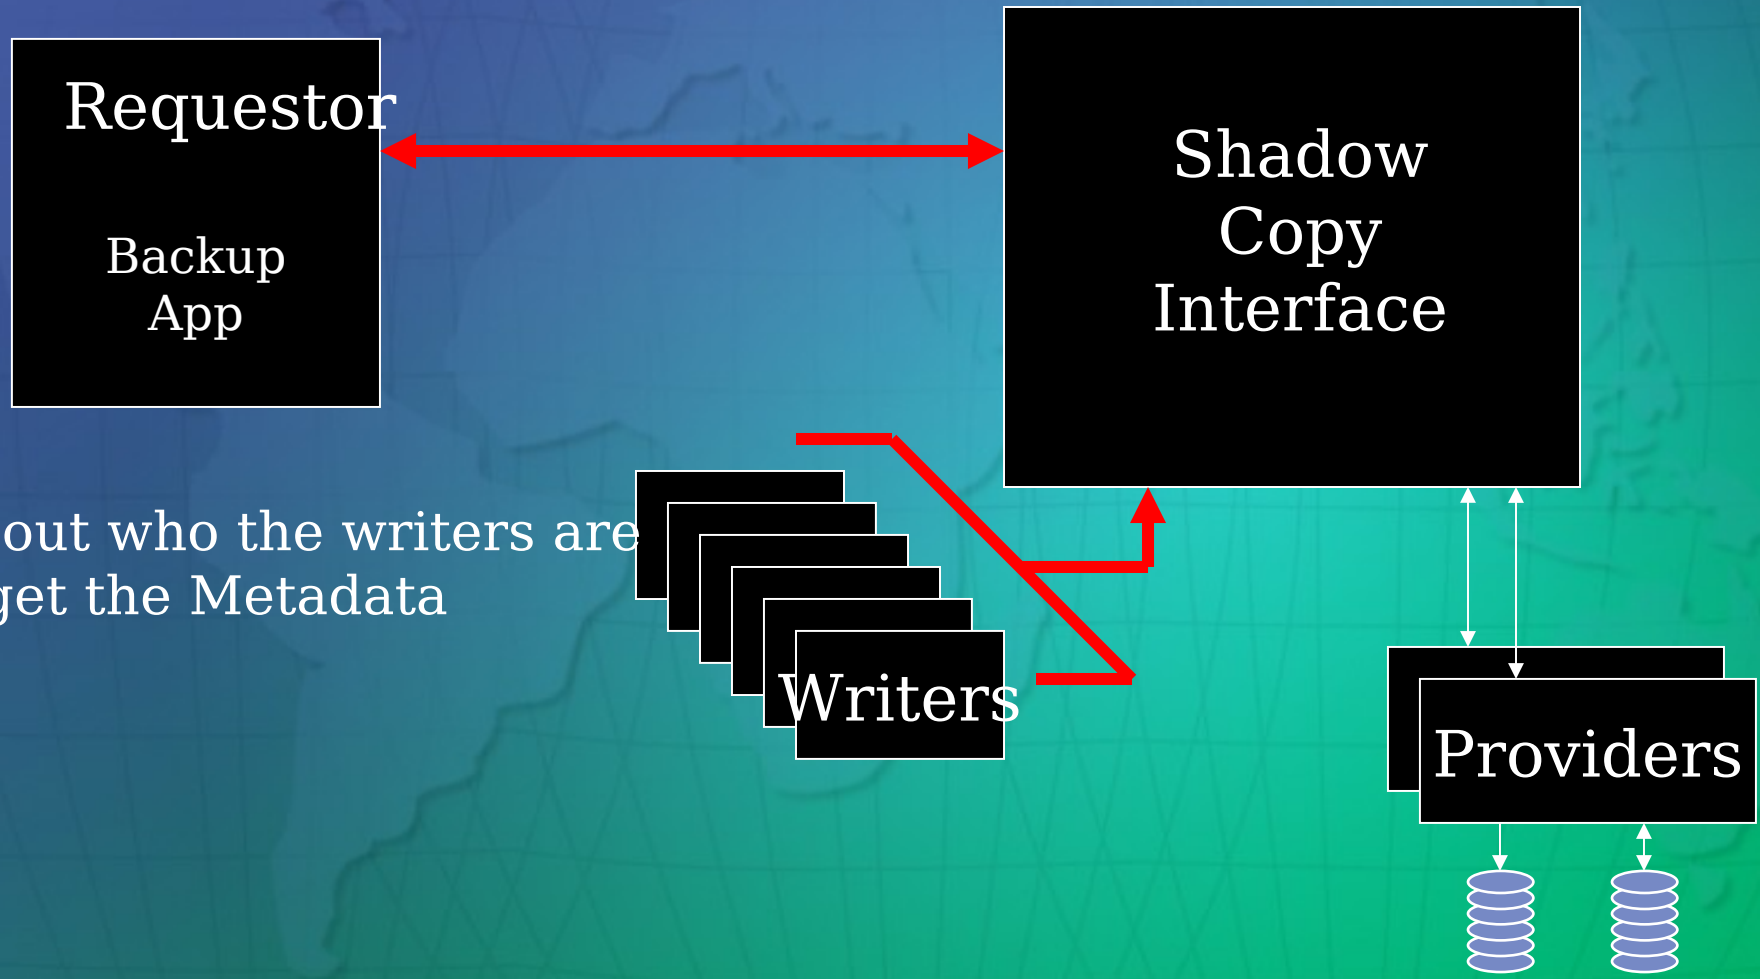
# Backup With Shadow Copies

## Block Diagram

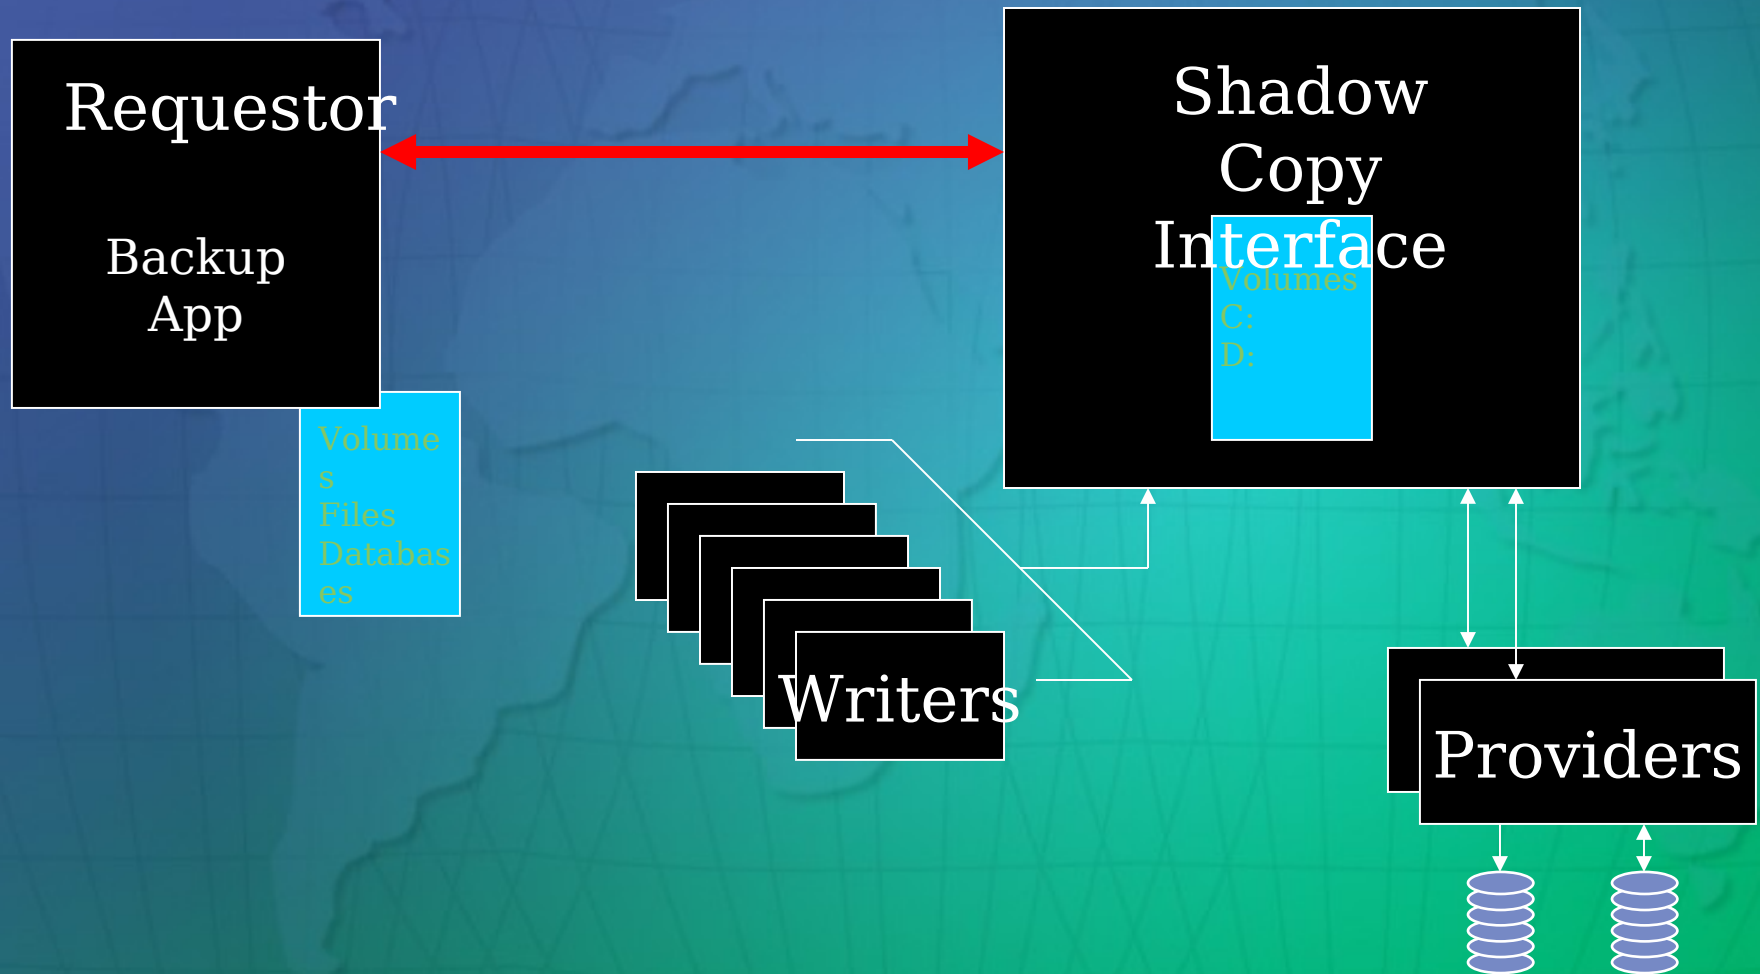




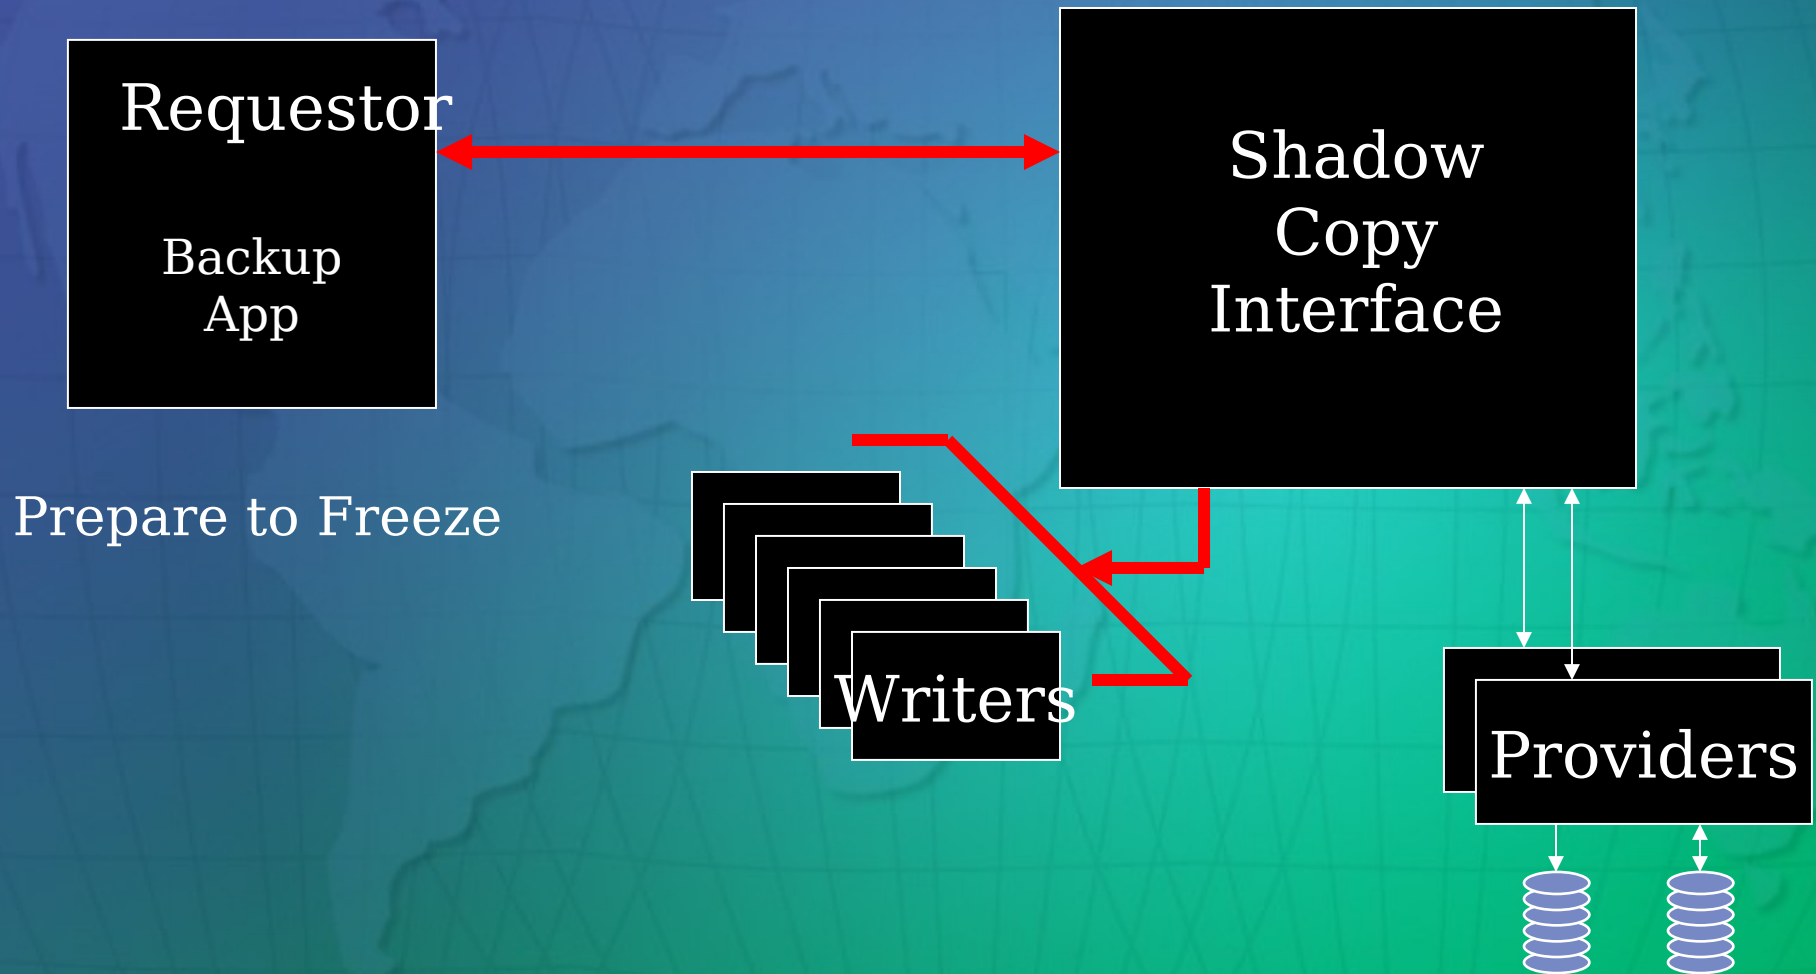
# Enumerate The Writers And Gather The Writer Metadata



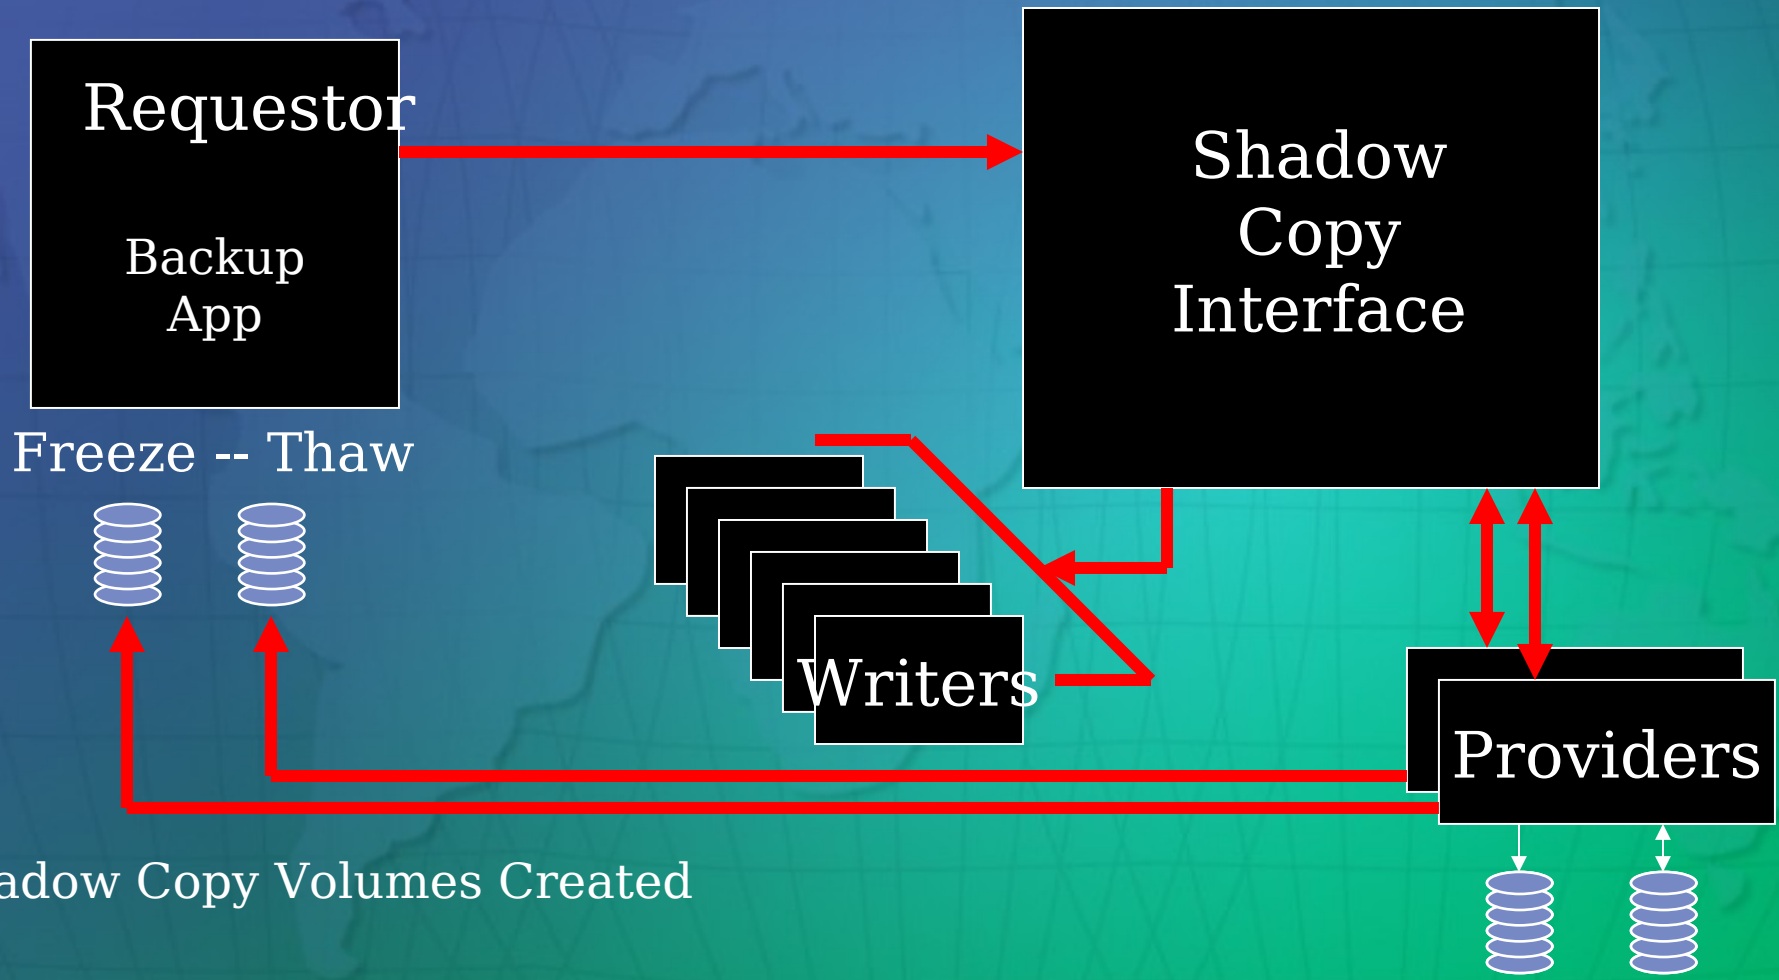
# Define The Shadow Copy Set



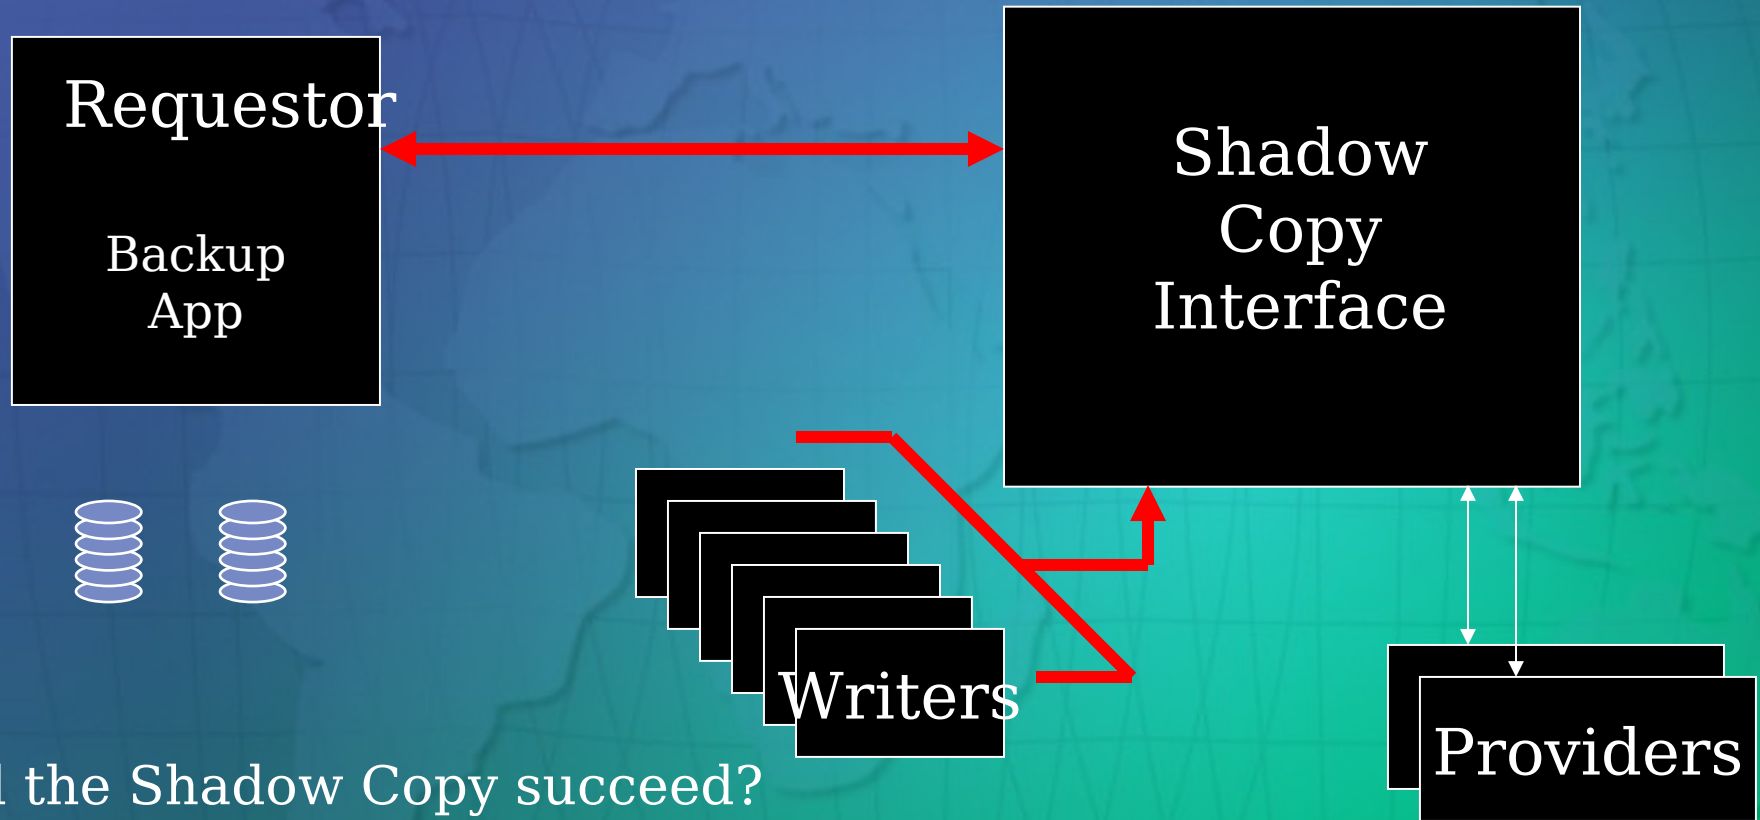
# Notify The Writers That The Shadow Copy Is Commencing



## Initiate Shadow Copy

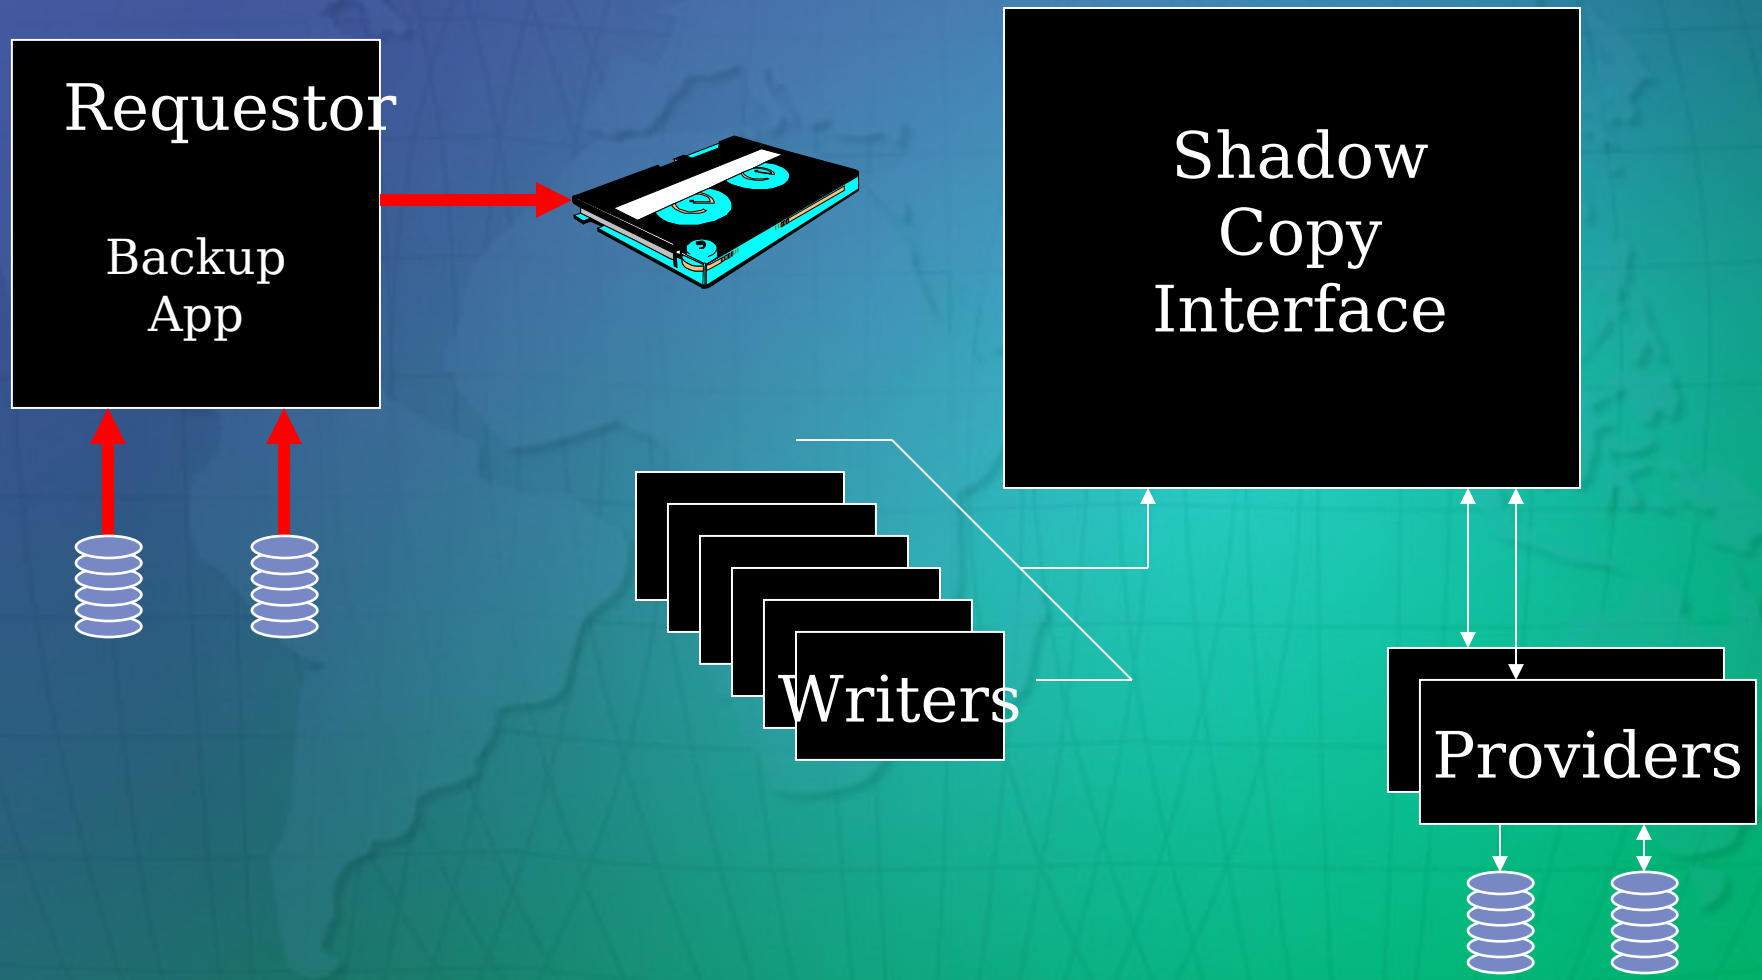


## Get Status From The Writers

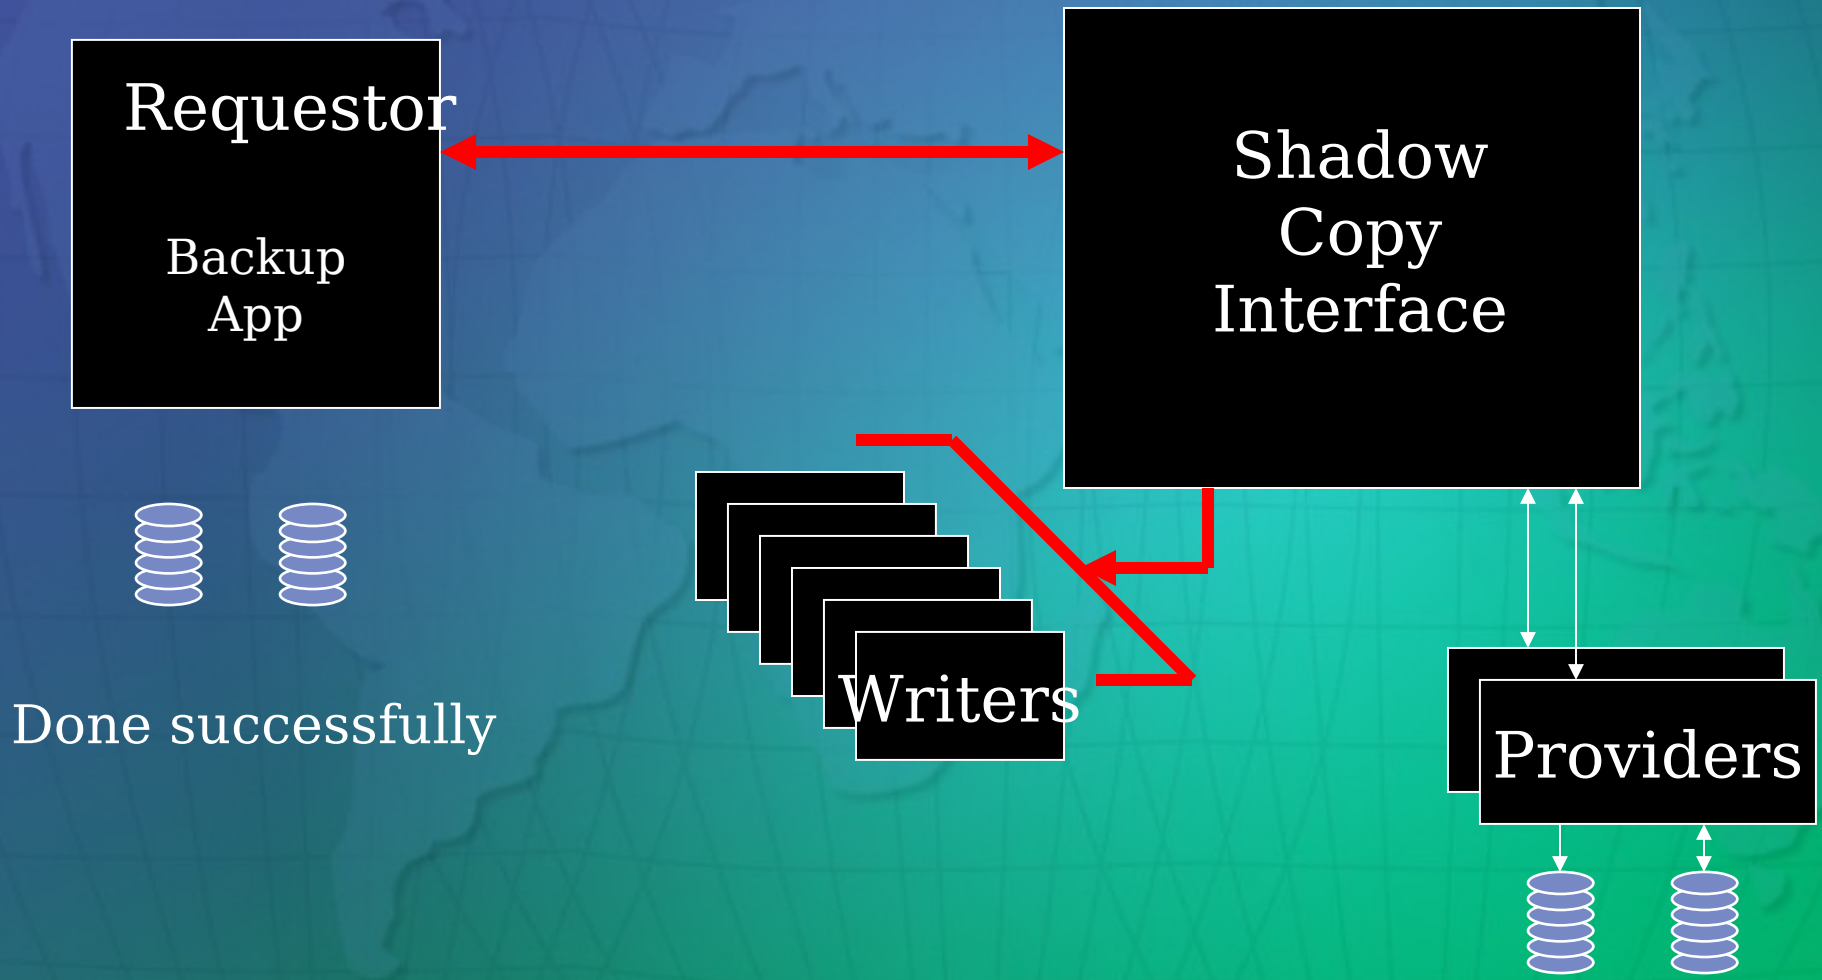




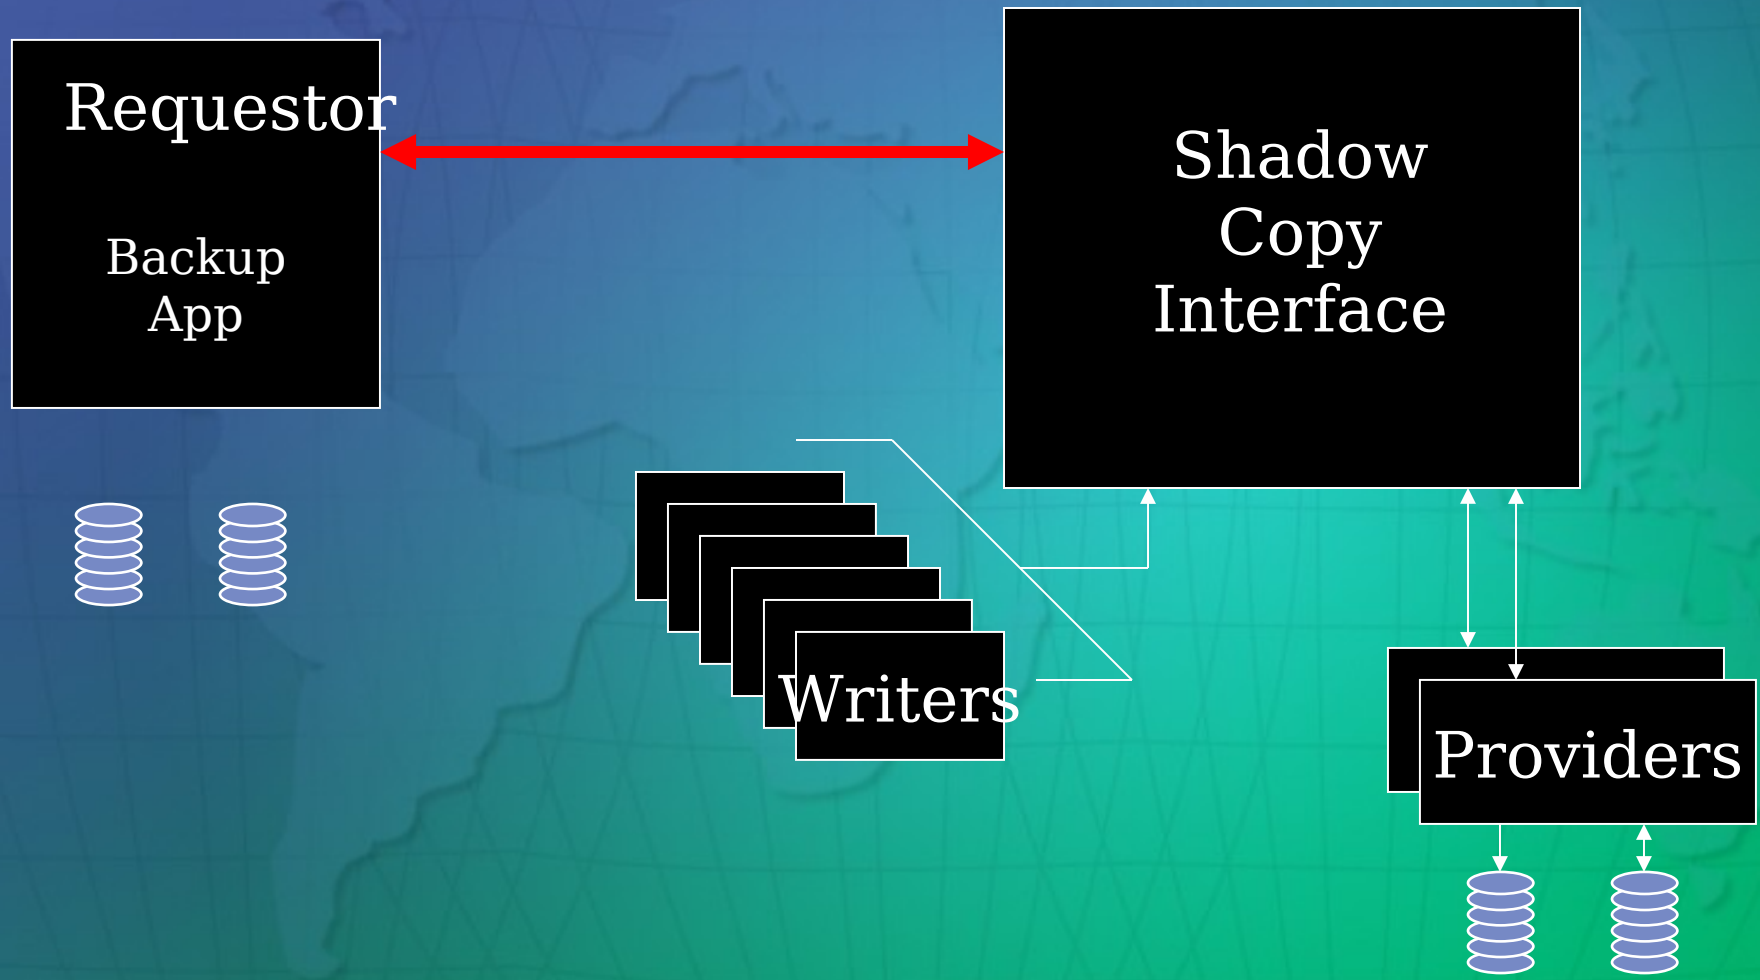
# DO THE BACKUP



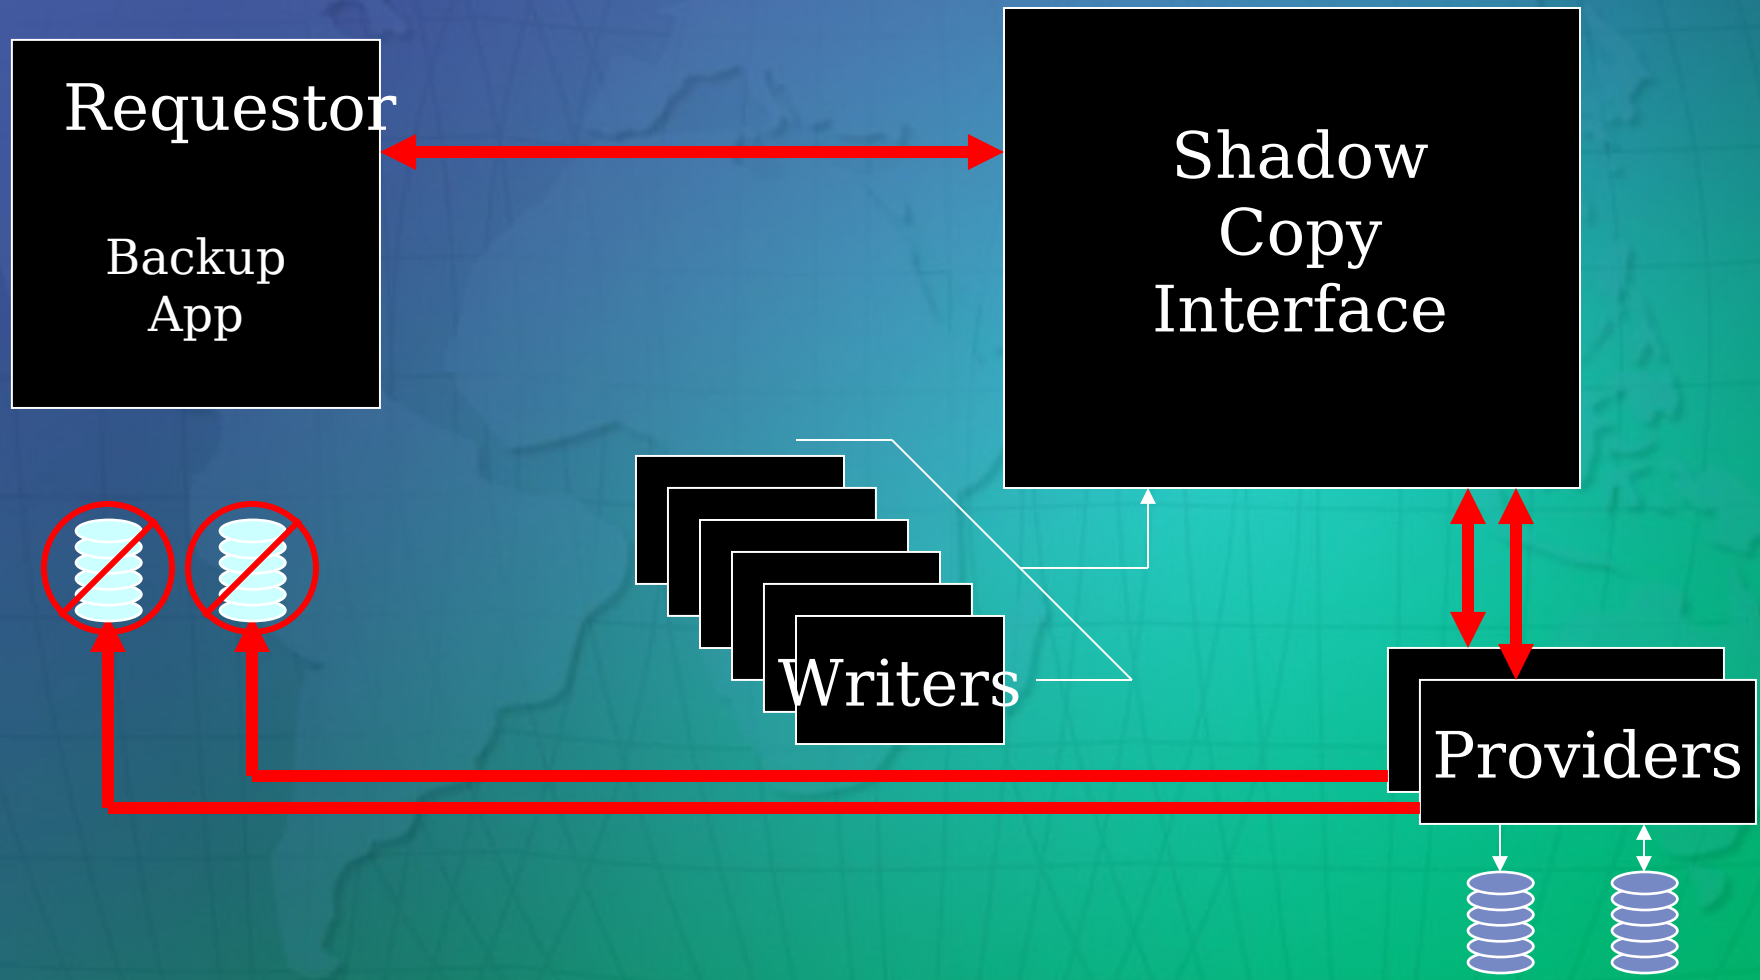
## Indicate Success To Writers



# Backup Is Completed



# Destroy The Shadow Copy

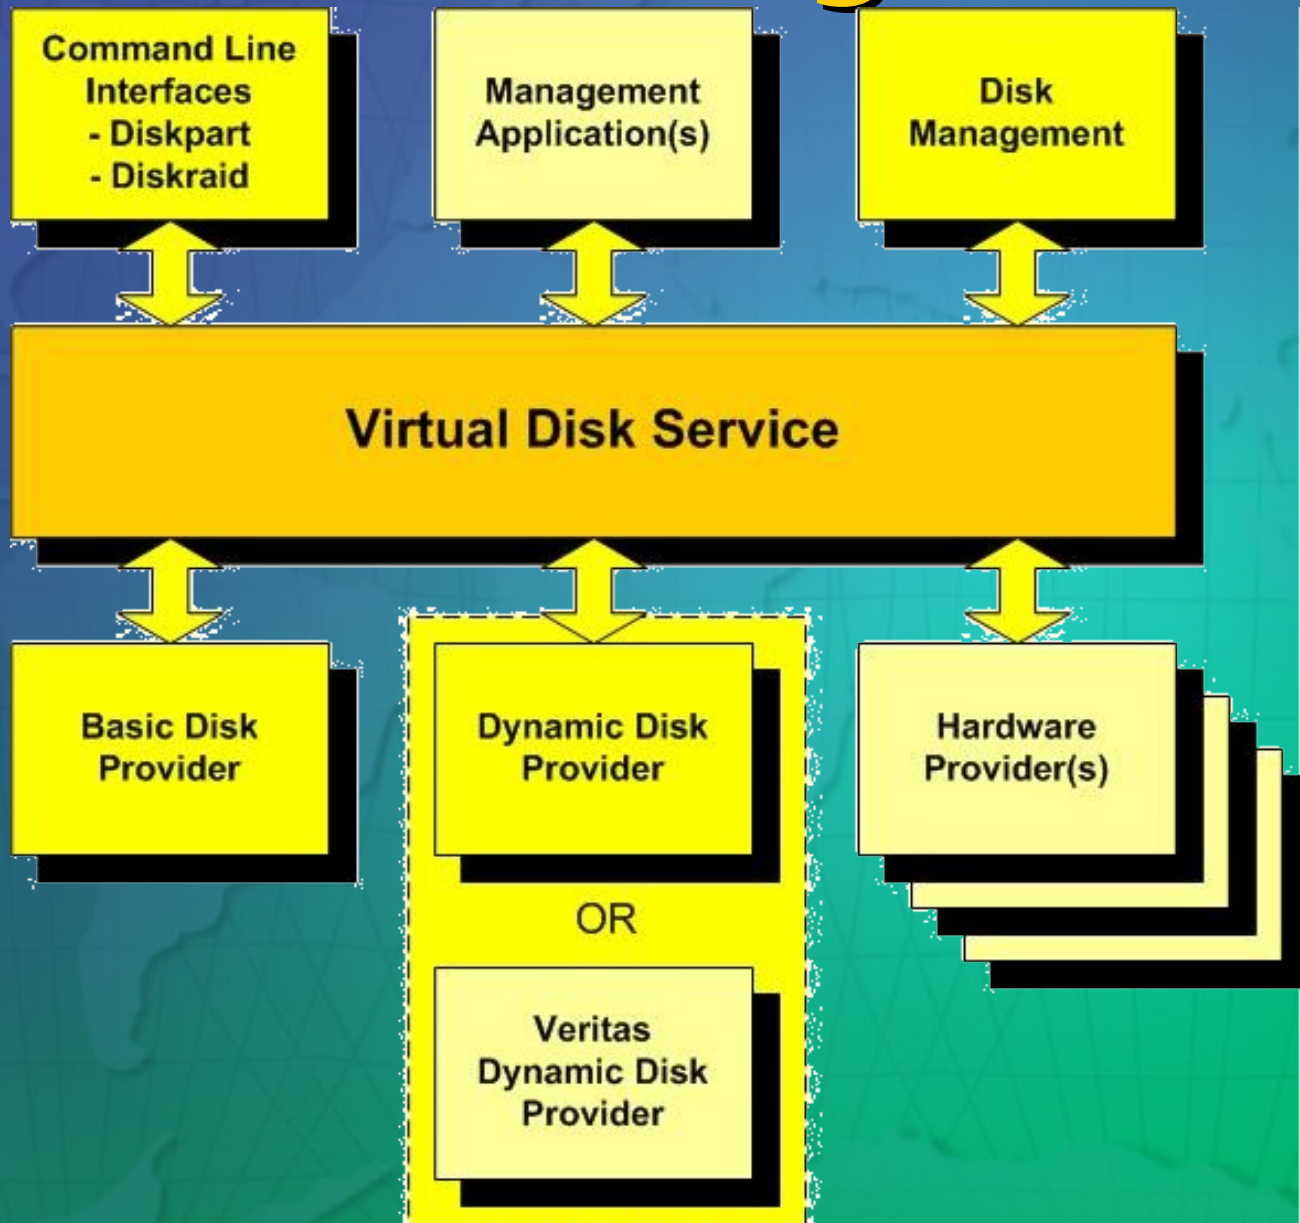


# **Virtual Disk Service (VDS)**

- **Vendor neutral API for RAID and block storage virtualization**
- **Benefits:**
  - **Introduces a common command line for hardware and software RAID**
  - **API in future versions**
- **Goals:**
  - **Vendor & Technology neutral**
  - **Manage logical volumes**
  - **Manage LUNs**



# VDS Block Diagram



# **.NET VDS Features**

- **VDS coordinator service**
- **Software provider for Basic disks**
- **Software provider for Dynamic disks**
- **DISKPART command line interface additional commands**
- **Disk Management (non-user visible) VDS implementation**
- **Hardware provider for PCI RAID using Storport**
- **DISKRAID CLI (available in Resource Kit)**


# **VDS Providers**

- **In-box providers**
  - **Basic disk**
  - **Dynamic disk**
  - **RAIDport**
  - **Windows XP/Windows 2000 down rev (remoting)**
- **Other providers**
  - **Veritas add-on**
  - **RAID IHVs**

# VDS Across Machines

- App runs anywhere
- Service runs everywhere
- Software provider runs wherever volumes are mounted
- IHV provider can run anywhere or only in one place

**Allows SAN install**





# **VDS Roadmap (After .NET)**

- **Software LUN masking**
- **Multipath configuration**
- **SAN device discovery**
- **Support for remote hardware provider**  
(resident in SAN)
- **Zoning/fabric management integration with Fabric Virtualization Service**
- **Two site or GeoMirroring**



# **Disaster Recovery (ASR)**

- **Disaster:**
  - **Physical destruction of the computer system (fire, earthquakes, etc.)**
  - **Catastrophic hardware failure (especially storage devices)**
- **Recovery:**
  - **Restore hardware configuration, base OS**
  - **Restore user data as a separate process**

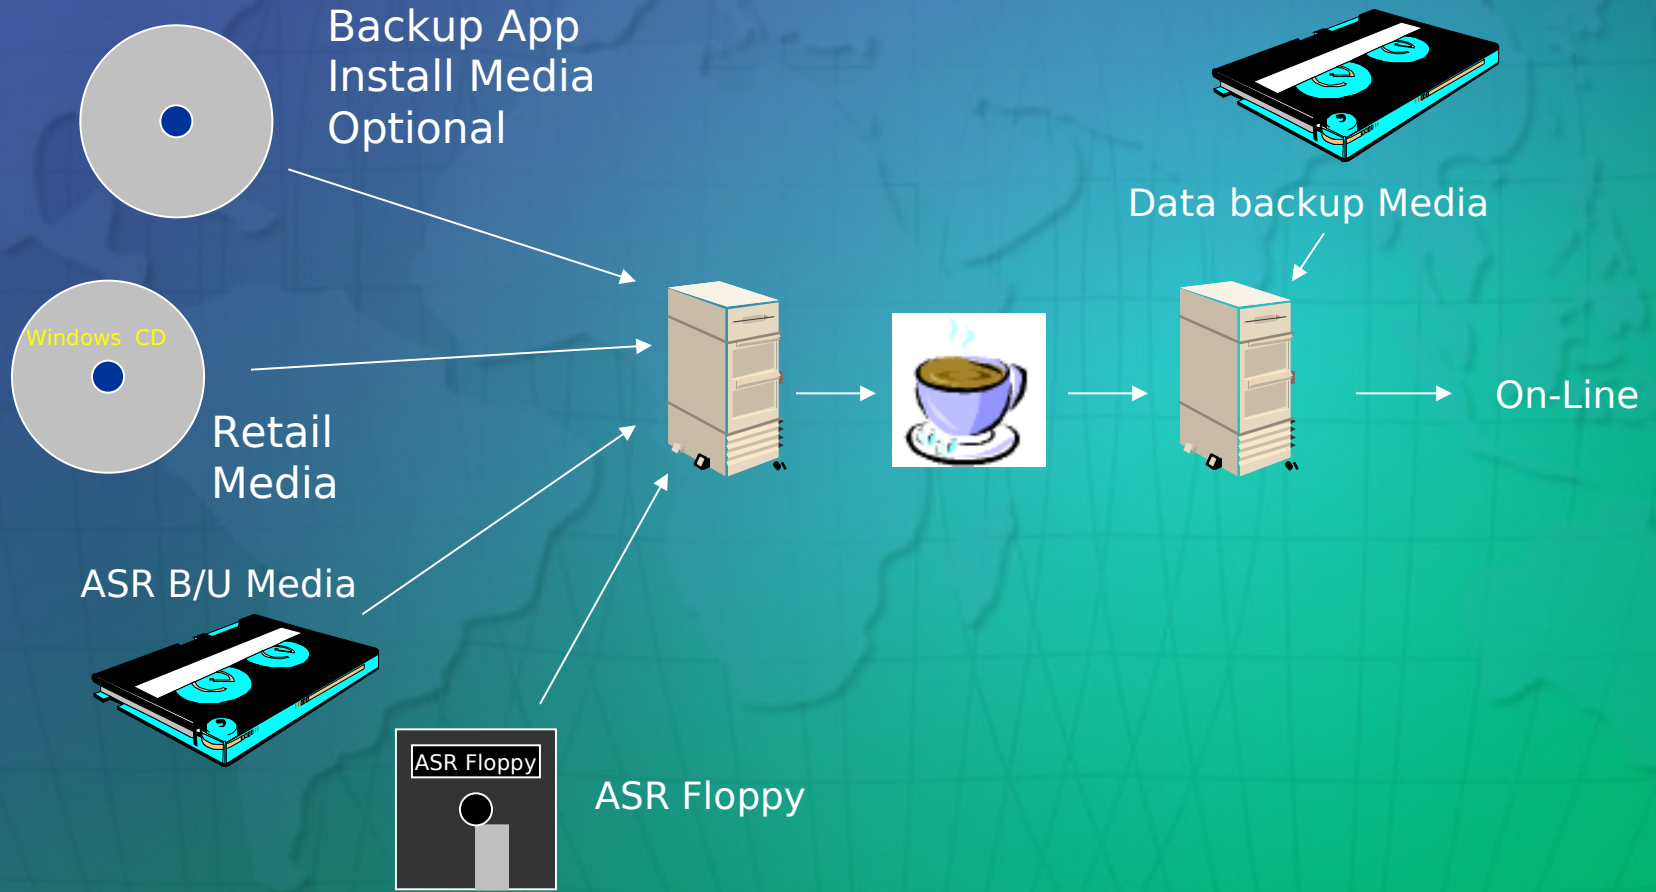
# **(old) Manual System Recovery**

- **User (administrator) gets new hardware**
- **User installs Windows**
- **User configures all physical storage to their original settings**
- **User installs backup-and-restore application, backup media drivers**
- **User recovers operating system**
- **User reboots and adjusts services**
- **User restores data**

# **Automated System Recovery**

- **New in Windows .NET Server and XP**
- **User (administrator) gets new hardware**
- **From the Windows CD, user runs ASR**
- **Inserts other media when prompted**
- **User restores data**

# What Is ASR



# ASR Goals

- **To bring a non-bootable system to a state from which a backup-and-restore application can be executed**
  - **To re-configure physical storage to its original state**
  - **To restore the operating system, applications and all settings**
- **To provide a mechanism for third-party vendors to incorporate ASR-related features**



# Target System Requirements

- Hardware must be identical as original system, except hard disks, video cards and NIC's
- Enough disks to restore all the Critical System Disks
- Storage capacity of each critical disk must be  $\geq$  corresponding original disk. Disk geometries must be compatible
- ASR state file (asr.sif) must be accessible through a local floppy drive

# ASR And Clusters

- **Backup:**
  - On each node, stores configuration info about local and all shared disks
- **Restore:**
  - While restoring the first node, all shared disks are restored (Lone-wolf up and running)
  - For other nodes, only local critical disks will be restored

# **ASR Versus Emergency Repair**

- **ER: Gone**
  - **Replaced missing or corrupt system files**
  - **Did not format drives or reconfigure storage**
- **ASR: System has been through a disaster**
  - **Boot partition is ALWAYS formatted, system partition may be formatted**

# **ASR Supported Configurations**

- **Basic and Dynamic Disks**
- **x86 and ia64 platforms**
- **MBR and GPT (EFI) disks**
- **Need a floppy drive**
- **ASR restore not networkable/remotable**

# **ASR With RIS**

- **The ASR process can be completely automated with the use of a RIS server**
  - **No floppy required**
  - **No F2 to start ASR**
  - **Uses PxE boot**
- **Whitepaper in preparation...**



# Command Line Scripting

- **FSutil**
- **Diskpart**
- **DiskRAID**
- **VSSAdmin**

# FSutil

- **Quota**      **Quota management**
- **Behavior**      **Control file system behavior**
- **Dirty**      **Manage volume dirty bit**
- **Volume**      **Volume management**
- **File**      **File specific commands**
- **Fsinfo**      **File system information**
- **Hardlink**      **Hardlink management**
  
- **Objectid**      **Object ID management**
- **reparsepoint management**      **Reparse point**
- **Sparse**      **Sparse file control**
- **Usn**      **USN management**

# Diskpart

- **ADD** - Add a mirror to a simple volume.
- **ACTIVE** - Activates the current basic partition.
- **ASSIGN volume.** - Assign a drive letter or mount point to the selected
- **BREAK** - Break a mirror set.
- **CLEAN off the** - Clear the configuration information, or all information,
- **disk.**
- **CONVERT** - Converts between different disk formats.
- **CREATE** - Create a volume or partition.
- **DELETE** - Delete an object.
- **DETAIL** - Provide details about an object.
- **EXIT** - Exit DiskPart
- **EXTEND** - Extend a volume.
- **HELP** - Prints a list of commands.
- **IMPORT** - Imports a disk group.
- **LIST** - Prints out a list of objects.
- **ONLINE** - Online a disk that is currently marked as offline.
- **REM** - Does nothing. Used to comment scripts.
- **REMOVE** - Remove a drive letter or mount point assignment.
- **RESCAN** - Rescan the computer looking for disks and volumes.
- **RETAIN** - Place a retainer partition under a simple volume.
- **SELECT** - Move the focus to an object.

# DiskRAID

- **Command line to control VDS**

# VssAdmin

- Create and delete shadow copies
- List shadow copies, writers and providers
- Manage diff area



# **Command Line Utilities**

demo

# High Availability Considerations

- Cluster configuration stored as part of system state for ASR-enabled backups
  - Allows restore of a single cluster node
  - Complete restore of a cluster configuration
  - In addition we have a new reskit tool that enables easier recovery after loss of a shared diskv(ClusterRecovery for Windows 2000 and .NET server)
- Partitions on basic disks can be extended online if a LUN is extended (by storage controller) using DiskPart
  - For both standalone and clustered servers

# Considerations For A Highly Available File Service

- **Different techniques for a highly available file service**
  - DFS with replicated root and content
  - Failover clustering
    - Simple file shares
    - Failover DFS roots and file shares
    - Replicated roots combined with failover file shares
- **Need to consider ways to make data highly available**
  - Multiple paths to data (multiple independent fiber channel paths)
  - Redundant storage controllers
  - Different RAID levels
    - Mirroring (RAID-1)
    - Mirroring + Striping (RAID 10; RAID 1 + 0)
    - RAID-5

# Server Cluster New Features

## Highlights for file server deployments

- 1 of 2**
- Cluster Configuration Wizard
    - Simplified cluster setup and configuration
  - Active Directory® Integration
    - End-to-end Kerberos authentication and delegation
  - Larger Clusters
    - 4 nodes in Advanced Server
    - 8 nodes in Datacenter
  - Automatic Server Recovery (ASR) support
    - Restores disk signatures for all disk (including shared disks)
    - Rebuilds a single node including cluster registry state
    - Can be used to rebuild a complete cluster (including quorum disk)
  - Rolling upgrade Windows NT® 4.0 to Windows 2000 to Windows .Net Server
  - WMI support for management and events



# Server Cluster New Features

Cluster account password utility  
Change cluster service account password without taking down the cluster

## Highlights for file server deployments

2 of 2

- **Enhanced SAN support**
  - Targeted reset of storage devices rather than full bus reset - requires supported by HBA device drivers
  - “Single storage bus” support
- **File system enhancements**
  - Physical disk resources no longer require drive letters
  - Mountpoints supported in a cluster (support > 23 disks)
  - Client-side caching for clustered SMB shares
  - Multiple DFS roots in a cluster
- **Whitepaper available with complete list of new features**



# Server Cluster Topologies

- 4 nodes in Advanced Server
- 8 nodes in Datacenter
- Fully functional cluster support for ia64

<b>Operating System</b>	<b># of Nodes</b>	<b>Storage Interconnect</b>
<b>Advanced Server (x86)</b>	<b>2</b>	<b>SCSI or Fibre</b>
<b>Advanced Server (x86)</b>	<b>&gt; 2 nodes</b>	<b>Fibre Channel</b>
<b>Datacenter (x86)</b>	<b>2 - 8</b>	<b>Fibre Channel</b>
<b>Advanced Server (ia64)</b>	<b>2 - 4</b>	<b>Fibre Channel</b>

# ***Microsoft***<sup>®</sup>